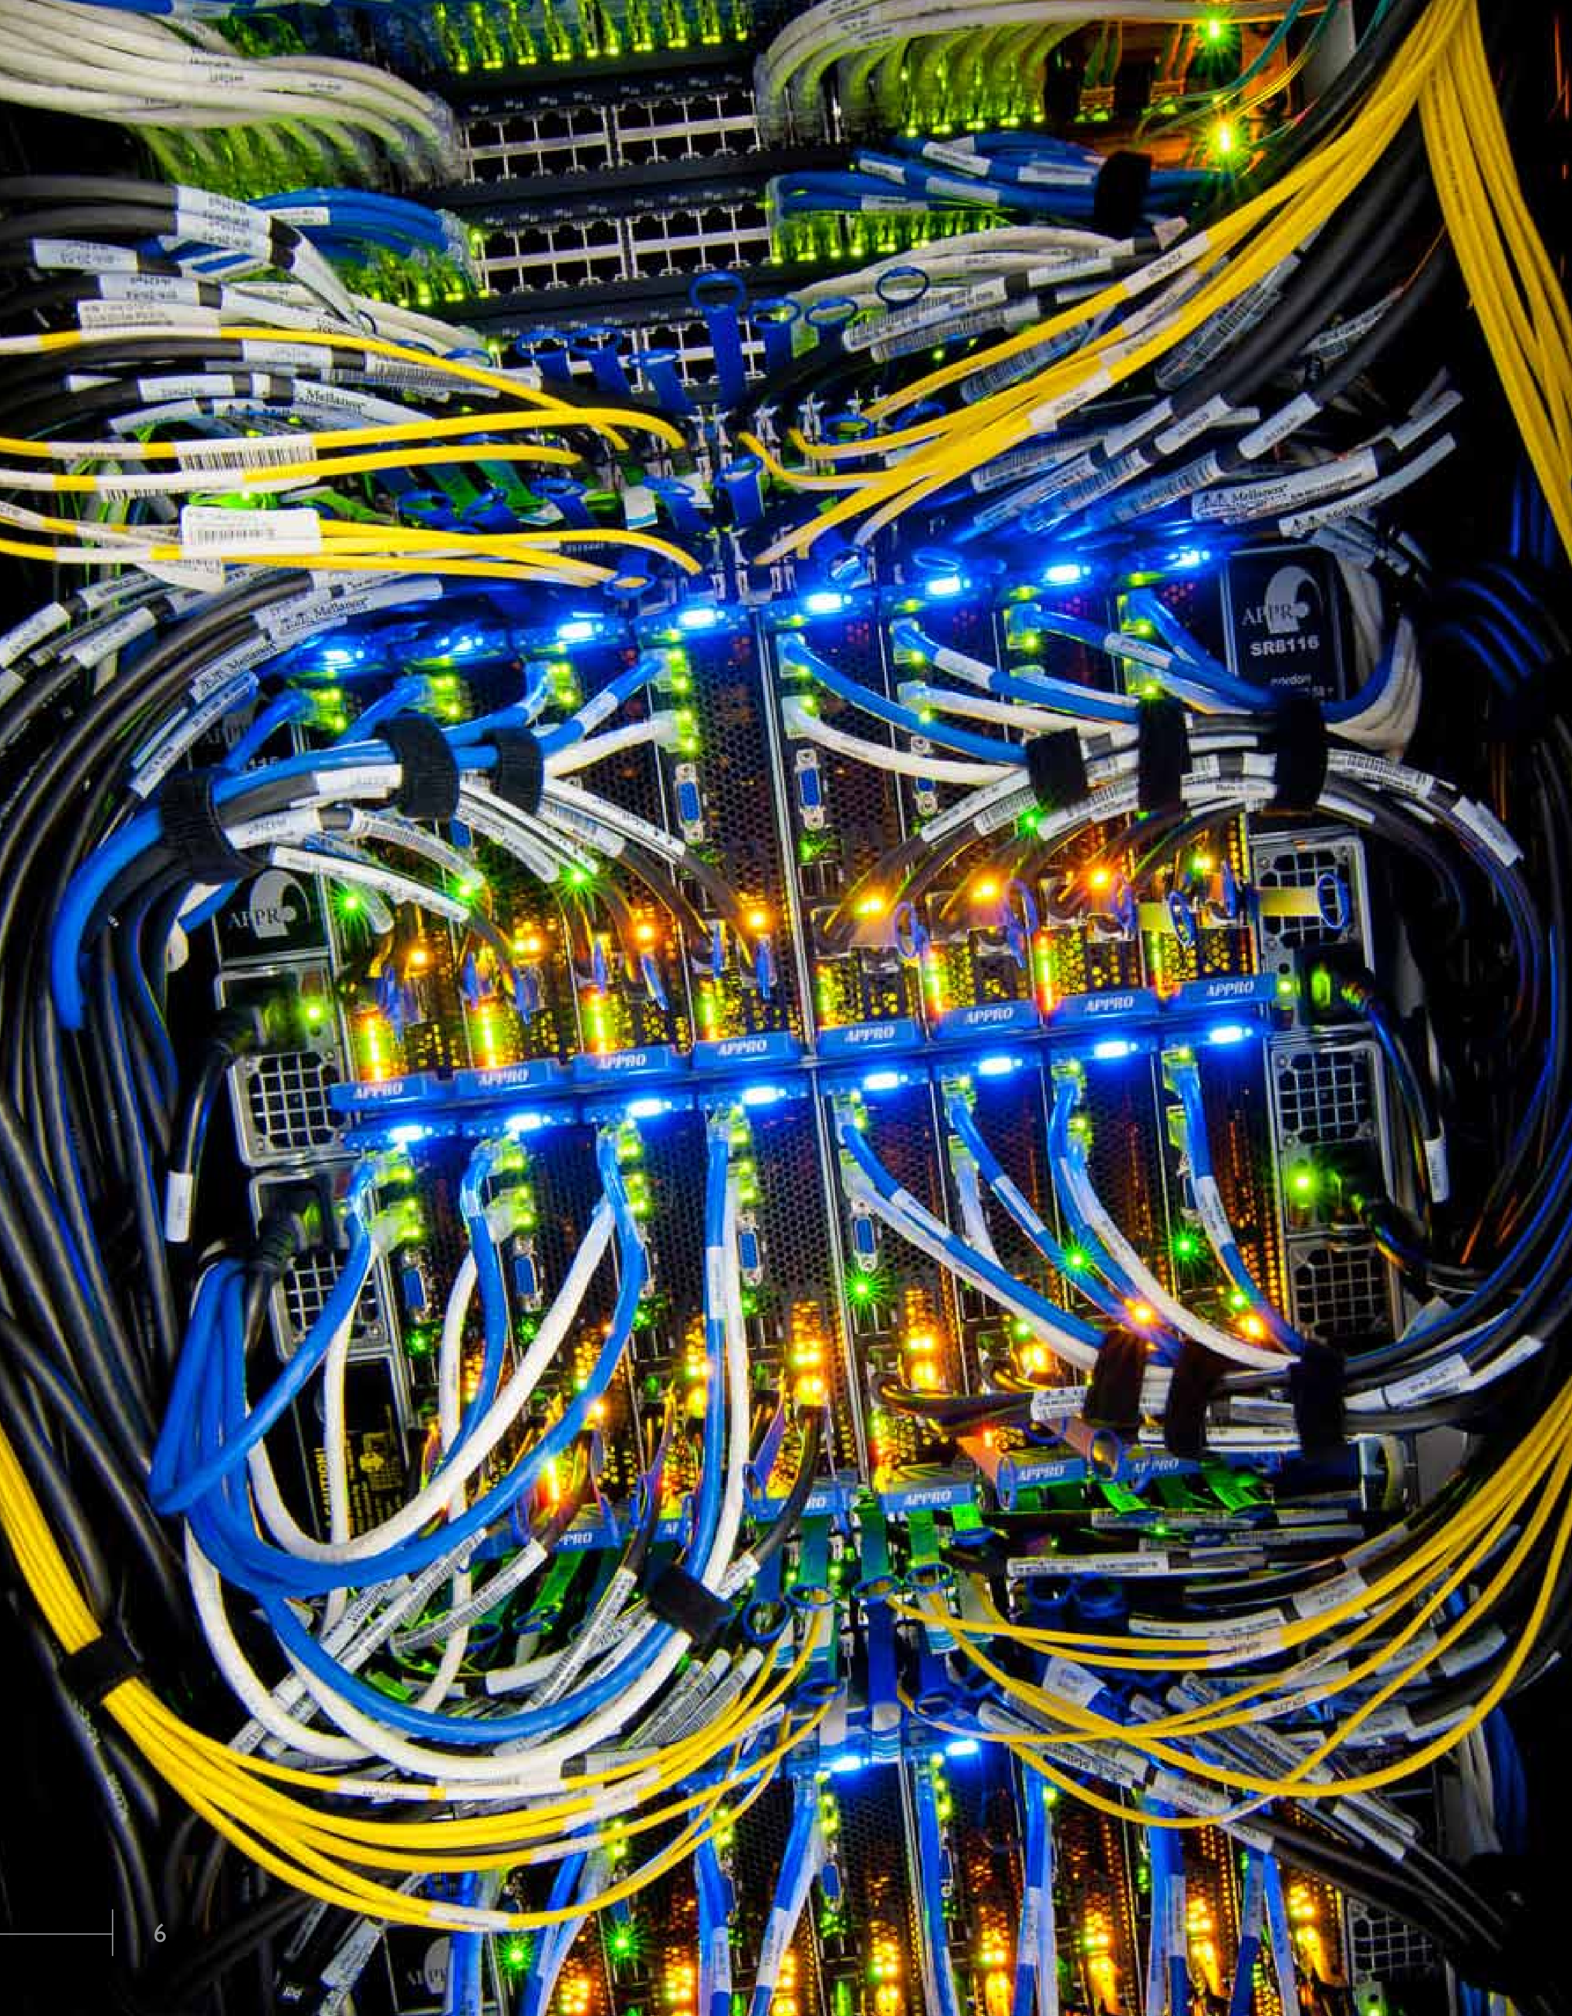
APPRO
SR8116

SDSC's BIG DATA CAPABILITIES

For almost 30 years SDSC has strived to provide a unique blend of advanced computing and expertise that serves as a local, state, and national gateway for collaboration, innovation, and education for an era defined and transformed by its information. Guided by this vision, SDSC not only has provided the resources and expertise needed for current research, the Center has sought to anticipate researchers' future requirements through the development and creative use of new tools to advance discovery.

Today, scientific research has entered what is called the "fourth paradigm": data-enabled investigation. That data comes in many forms and from all over the globe—academic and government institutions, commercial enterprise, even observational tools such as satellites and deep-sea sensors. It is streaming in at ever-increasing rates, creating a formidable challenge for researchers to not only extract meaningful value from this data, but to manage, organize, and store massive amounts of information.

SDSC is meeting this multi-faceted challenge by focusing on three key areas of capability: High-Performance Computing and Resources (for rapid data processing); Comprehensive Data Management (for secure and stable storage); and Networking/Connectivity (for rapid data transmission).

Inside SDSC's *Gordon* supercomputer. Image: Erik Jepsen, UC San Diego (left)

SDSC's first supercomputer, the Cray XMP-48, in 1985 (right)

Trtion Shared Compute Cluster or TSCC servers

HIGH-PERFORMANCE COMPUTING AND RESOURCES

Since its inception in 1985, SDSC has been at the forefront of high-performance computing, developing innovations to make supercomputers more capable, compact, and energy-efficient by orders of magnitude. In addition to smaller clusters dedicated to specific tasks, SDSC currently has three main HPC systems used by researchers at the local, national, and global levels across numerous domains that include both traditional and non-traditional areas when it comes to using advanced computation.

Trestles

High-Productivity Workhorse

Trestles came online in early 2011 to provide researchers not only significant computing capabilities, but to allow them to be more computationally productive. Known throughout the national HPC community as a high-productivity workhorse, *Trestles* is based on the idea that by tailoring a system for the majority of jobs rather than a handful of researchers who run jobs at thousands of core counts, users would be rewarded with high throughput and increased scientific productivity. Today, *Trestles* is recognized as the leading science gateway platform in the National Science Foundation's eXtreme Science and Engineering Discovery Environment (XSEDE) portfolio, with more than 650 users per month run through the popular CIPRES phylogenetics portal alone. *Trestles* users span a wide range of domains, including astronomy, biophysics, climate sciences, computational chemistry, material sciences, and more.

Trestles will be replaced in 2015 by an all-new HPC system called *Comet*, a petascale supercomputer designed to transform advanced scientific computing by expanding access and capacity among traditional as well as non-traditional research domains.

Gordon

Delivering on Data-intensive Demands

Gordon entered production in early 2012 as one of the 50 fastest supercomputers in the world—and the first one to employ massive amounts of flash-based memory, making it hundreds of times faster than conventional HPC systems while having enough bandwidth to handle extremely large datasets. The result of a five-year, $20 million NSF grant, *Gordon* has 300 trillion bytes of flash memory and 64 I/O nodes, making the system ideal for researchers who need to sift through tremendous amounts of data. In effect, *Gordon* is designed to do for scientific research what Google does for web searches.

By the end of 2013, 799 research projects using *Gordon* were awarded among 479 principal investigators. In early 2013, *Gordon* completed its most data-intensive task so far: rapidly processing raw data from almost one billion particle collisions as part of a project to help define the future research agenda for the Large Hadron Collider (LHC). Under a partnership between a team of UC San Diego physicists and the Open Science Grid, *Gordon* provided auxiliary computing capacity by processing massive data sets generated by one of the LHC's two large general-purpose particle detectors used to find the elusive Higgs particle. The around-the-clock data processing run on *Gordon* was completed in about four weeks' time, making the data available for analysis several months ahead of schedule (see page 20).

Gordon, housed in SDSC's Datacenter, is the first supercomputer to employ vast amounts of flash-based memory. Image: Erik Jepsen, UC San Diego.

Rick Wagner is SDSC's High-Performance Computing Systems Manager, responsible for SDSC's Linux-based computing clusters and related systems that support users across a broad range of scientific disciplines.

In addition to leading-edge genomic research (*Gordon* is capable of storing 100,000 entire human genomes) the system is now being used by Arieah Warshel, who won the 2013 Nobel Prize in chemistry for developing detailed computer simulations of complex chemical processes for tasks such as designing new drugs or solar cells. Warshel is using *Gordon* and other SDSC data storage resources exclusively after being allocated 3 million CPU (core processing unit) hours last year.

Triton Shared Computing Cluster (TSCC)

Affordable Computing for Campus and Corporate Needs

In mid-2013, SDSC deployed a new high-performance research computing system called the *Triton Shared Computing Cluster*, or *TSCC*, to serve researchers at UC San Diego and other UC campuses, as well as external academic, non-profit, and corporate users. *TSCC* is operated by SDSC for UC San Diego's Research Cyberinfrastructure (RCI) program. The cluster features state-of-the-art hardware, and a revamped participation model that provides researchers more options for funding their computing-based needs. Primary benefits to participants include gaining access to a much larger resource than they could afford solely for their labs, and having a system that is professionally maintained by full-time staff at SDSC.

Enabled with a wide range of open source and licensed software for science and engineering, *TSCC* is designed to work with the full complement of capabilities at UC San Diego, including high-performance networking and centralized storage systems. Connectivity is available to high-throughput scientific instruments across campus, such as DNA sequencers and mass spectrometers.

Phil Papadopoulos is SDSC's Chief Technology Officer and chief architect behind SDSC's *Data Oasis* storage system. Papadopoulos also is principal investigator for the Prism@UCSD project to build a campus cyberinfrastructure capable of supporting extreme data-intensive communications.

COMPREHENSIVE DATA MANAGEMENT

Developing and operating fast and robust HPC systems is only part of the equation. The ability to provide a secure and stable environment to store data, both for the short and long term, is essential for researchers both before and after their computational research is done. As with HPC interfaces, another requirement is that data storage and retrieval is user-friendly.

Data Oasis

Among Academia's Fastest Parallel File Systems

In 2012, SDSC "supercharged" *Data Oasis*, a Lustre-based parallel file storage system linked to *Trestles*, *TSCC*, and *Gordon*. As a critical component of SDSC's Big Data initiatives, *Data Oasis* currently has four petabytes of capacity and speeds of up to 100 gigabytes per second to handle just about any data-intensive project. Using the I/O power of *Gordon* and *Trestles*, those sustained transfer rates make *Data Oasis* one the fastest parallel file systems in the academic community. The sustained speeds mean researchers could retrieve or store 64 terabytes (TB) of data—the equivalent of *Gordon's* entire DRAM memory—in about 10 minutes, significantly reducing time needed for retrieving, analyzing, storing, or sharing extremely large datasets.

Big Data is not just about sheer size, but also about the speed of moving data where it needs to be, and the integrated infrastructure and software tools needed to effectively do research using those data. The capability of *Data Oasis* allows researchers to analyze data at a much faster rate than other systems, which in turn helps extract knowledge and discovery from these datasets.

SDSC CLOUD
Version 2.0

The *SDSC Cloud*, which went into production in late 2011, is the first large-scale academic deployment of cloud storage in the world. *SDSC Cloud* storage currently has 2.7 petabytes of raw space used by almost 400 partners. The system is being expanded from an object-based file store running on OpenStack's SWIFT platform to include other capabilities. *SDSC Cloud 2.0* includes Nova, OpenStack's cloud compute capability. This provides the opportunity to build Data as a Service (DaaS) on top of an open source, cloud-based technology, marrying SDSC's service offerings with its growing expertise in cloud computing and solid data management and engineering research. The platform was designed around the use cases of researchers, which includes SDSC's PIs working on data management and Big Data projects as well as campus, system, and collaborators. In 2014, SDSC plans to launch the PACE Starter Kit, a collection of cloud-based and virtual technologies to extend data mining sandboxes and environments following the popular Data Mining Boot Camps held by SDSC's Predictive Analytics Center of Excellence (PACE).

Christine Kirkpatrick is division director of SDSC's IT Systems and Services Division, which designs, deploys, and operates high-performance systems and provides services supporting a full range of academic and industry researchers.

NETWORKING/CONNECTIVITY

While rapid computational processing and stable storage structures are essential to conduct data-enabled science, the cyberinfrastructure must also include networks that allow fast and unrestricted flow of information between systems and researchers. The answer is using dedicated 'fat' pipes that can accommodate extreme-sized bursts of data. SDSC has helped lead the way on several fronts in this area.

Prism@UCSD

The HOV Lane for Broad-bandwidth Research

Working with campus partners, SDSC helped establish a research-defined, end-to-end networking cyberinfrastructure for the UC San Diego campus that is capable of supporting large data transmissions between facilities that might otherwise hobble the main campus network. Called Prism@UCSD and backed by a $500,000 NSF grant, researchers with the campus' California Institute for Telecommunications and Information Technology (Calit2) and SDSC began work on the network in 2013 to support research in data-intensive areas such as genomic sequencing, climate science, electron microscopy, oceanography, and physics. Slated for completion in late 2014, the project is already serving researchers with full functionality.

"One can think of Prism as the HOV lane, whereas our very capable campus network represents the other lanes on the freeway," says Philip Papadopoulos, principal investigator on the Prism@UCSD project and SDSC's chief technology officer.

Almost 150 communication fibers tunnel into UC San Diego's Atkinson Hall, including the main Calit2 server room (pictured above). From there, researchers have access to the main campus networks, as well as regional, national, and international wireless and optical networking test beds and visualization centers.

CHERuB

Connecting to the Information Superhighway

In late 2013 SDSC and UC San Diego's Administrative Computing and Telecommunications (ACT) organization were awarded a second $500,000 NSF grant to connect the campus to high-bandwidth national research networks to advance a new range of data-driven research. Called CHERuB for Configurable, High-speed, Extensible Research Bandwidth, the project will provide 100 gigabit-per-second connectivity—the new high-end for wide-area research networks. CHERuB will support multi-institutional data transit over networks such as the Internet2's Advanced Layer 2 Service (AL2S) and ESnet, as well as a joint project between those networks called the Advanced Networking Initiative (ANI), the result of a $62 million grant under the American Recovery and Reinvestment Act to build a national 100G "information backbone."

When completed, the CHERuB link will place UC San Diego among research universities and institutions having the highest available connectivity, with a capacity 10 times greater than existing modern data networks. CHERuB is the missing piece that will connect UC San Diego's Prism network to even faster national networks to advance scientific research.

Examples of research domains that will benefit from CHERuB include cosmology, atmospheric sciences, electron microscopy, genomic sequencing, oceanography, high-energy physics, and telemedicine—all of which can encompass data-rich research and whose advancements rely on multi-site or inter-institutional activities.